
Subject: Re: [PATCH 0/16] Pid namespaces
Posted by [Dave Hansen](#) on Mon, 09 Jul 2007 19:58:37 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Mon, 2007-07-09 at 09:58 +0400, Pavel Emelianov wrote:
> Dave Hansen wrote:
> > On Fri, 2007-07-06 at 12:01 +0400, Pavel Emelianov wrote:
> >> This is "submission for inclusion" of hierarchical, not kconfig
> >> configurable, zero overheaded ;) pid namespaces.
> >
> > Pavel, I'm a bit disappointed that you went ahead and sent this. I
> > thought that, perhaps, you might have brought up how displeased you were
> > with Suka's patches when we discussed them at OLS.
> >
> > Hold your horses there a bit. This has "little" overhead for the common
> > case, which is a single level of pid namespaces. That means that it is
> > quick to access the "global" pid which would be the one that the "host
> > container" sees. It also provides quick access to the pid which a
> > containerized task gets when the task itself calls getpid(). This quick
> > access is provided by storing the values directly in the task struct.
> >
> > However, when there is more than one level in the container hierarchy,
> > the optimization breaks down. A process which exists in a three-level
> > hierarchy has slow access to the middle level pid. Your approach stores
> > this information in a linked list, and surely *that* is going to have
>
> No. This approach stores numerical values in array. I have
> removed the lists at all.

Ahh. I was confused by the hlist in 'struct pid'. You are very correct.

Suka actually coded up something very, very similar to what you have. I just made him remove some of it so that the patches could be more easily reviewed. I figured we could add the fully dynamic allocation later, which you have already done. That part of your patches is remarkably similar.

-- Dave

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
